

Hypergraph conditions for the solvability of the ergodic equation for zero-sum games

Marianne Akian^{*1}, Stéphane Gaubert^{*2} and Antoine Hochart^{*†3}

Abstract—The ergodic equation is a basic tool in the study of mean-payoff stochastic games. Its solvability entails that the mean payoff is independent of the initial state. Moreover, optimal stationary strategies are readily obtained from its solution. In this paper, we give a general sufficient condition for the solvability of the ergodic equation, for a game with finite state space but arbitrary action spaces. This condition involves a pair of directed hypergraphs depending only on the “growth at infinity” of the Shapley operator of the game. This refines a recent result of the authors which only applied to games with bounded payments, as well as earlier nonlinear fixed point results for order preserving maps, involving graph conditions.

Index Terms—Zero-sum games, stochastic control, ergodic control, risk-sensitive control, nonlinear consensus, computational methods, directed hypergraphs.

I. INTRODUCTION

a) Motivation: A general issue in the study of stochastic control and zero-sum game problems is to give conditions which guarantee that the mean payoff per time unit is independent of the choice of the initial state. A basic tool to address this issue is the ergodic eigenproblem. The latter is a nonlinear equation, involving a function or vector called *bias* or *potential*, and a scalar called *ergodic constant*. When the ergodic equation has a solution, the ergodic constant gives the mean payoff for every initial state. The ergodic equation is often called the *average case optimality equation* in stochastic control, see [1], [2]. It may be thought of as a nonlinear extension of the *Poisson equation* arising in potential theory.

Conditions which guarantee the solvability of the ergodic equation are generally referred to as *ergodicity conditions*. Such conditions typically involve the transition probabilities. One of the oldest conditions of this kind was obtained by Bather [3], who showed that the ergodic equation has a solution for the class of Markov decision processes that are *communicating*, meaning that for every pair of states i to j , there is a strategy and a horizon k under which the probability to reach j from i is positive. An interest of Bather’s condition is its structural nature: as it only depends on the transition probabilities, it is insensitive to a perturbation of the payments. However, a limitation of Bather’s

communication condition is that it is tied to the one-player case: the communication condition is almost never satisfied for two-player problems (including zero-sum deterministic games).

The question of finding generalized communication or ergodicity conditions has emerged in a number of works dealing with control and games [4], [5], [6], [7], [8], but also discrete event systems [9], [10], mathematical economy [11], and Perron-Frobenius theory [12], [13]. A similar problem has also been studied in the infinite dimensional setting, when dealing with Hamilton-Jacobi type PDE: equations arising from deterministic optimal control problems were dealt with in [14]; the understanding of ergodicity conditions for PDE arising from differential games is a difficult question of current interest, see [15].

b) Main result: In the present paper, we address the question of the solvability of the ergodic equation for zero-sum game problems with *finite* state space, but with arbitrary action spaces. Our main result, Theorem 6, provides a general sufficient condition for the solvability of the ergodic equation, involving an accessibility condition in directed hypergraphs. These hypergraphs only depend on the “behavior at infinity” of the dynamic programming operator of the game. Intuitively, a hyperarc encodes the possibility for one of the two players to force the access in one step, with positive probability, to a given set of states, independently of the action of the other player. When the action spaces are finite, Theorem 6 has a simple game theoretic interpretation: it shows that the ergodic equation is solvable for all perturbations of the payment function if, and only if, we cannot find two disjoint subsets of states I, J such that Player MAX, starting from any state of J , has a possibility to force the next state to avoid being in I , with positive probability, whereas Player MIN, starting from any state of I , has a possibility to force the next state to avoid being in J , still with positive probability.

c) Related results: The present results refine or generalize several earlier ergodicity conditions. In particular, the present hypergraph condition (Theorem 6) refines the “generalized Perron-Frobenius theorem” of Gaubert and Gunawardena [13], which gave a sufficient condition for the solvability of the ergodic equation in terms of graph accessibility (hypergraphs lead to tighter conditions). It also extends results of Cavazos-Cadena and Hernández-Hernández [7], which apply to a subclass of operators arising from risk sensitive control problems. It finally refines a recent result of the authors [8], in which reachability conditions involving more special hypergraphs were introduced for games with

^{*}M. Akian, S. Gaubert and A. Hochart are with INRIA Saclay-Ile-de-France and CMAP Ecole polytechnique, Route de Saclay, 91128 Palaiseau Cedex, France

[†]A. Hochart is supported by a PhD fellowship of Fondation Mathématique Jacques Hadamard (FMJH)

¹marianne.akian@inria.fr

²stephane.gaubert@inria.fr

³antoine.hochart@cmaph.polytechnique.fr

bounded payments. The present Theorem 6 is more general, since arbitrary payments are allowed, and when the payments are bounded, the main result of [8] can be recovered as a special case of the present one, as the hypergraph conditions of Theorem 6 reduce to the ones in [8]. The case of *unbounded payments* is very common when the state space is not compact. Nevertheless, examples with unbounded payments and finite state space arise typically in the modeling of safety issues, where one player controls the transition probability to an undesirable state with a cost blowing up as this probability vanishes.

The proof of our main result further develops an idea of [13]: checking the boundedness of the so called “slice spaces” with respect to Hilbert’s seminorm. In this respect, we solve one question left open in [13] (characterizing combinatorially the boundedness of all slice spaces).

The paper is organized as follows. In Section II, we recall some basic definitions and results concerning zero-sum stochastic games. In Section III, we state and establish the main result (ergodicity condition, Theorem 6). This section includes a discussion of algorithmic and computational complexity aspects. In Section IV, we compare Theorem 6 and its corollary with earlier results, and provide an illustrative example.

II. PRELIMINARIES ON ZERO-SUM STOCHASTIC GAMES

A. The mean payoff

In this paper, we consider zero-sum stochastic games with *perfect information* and finite state space. We will denote by $S := \{1, \dots, n\}$ the state space, and by $\Delta(S)$ the set of probability measures on S . For each state $i \in S$, we denote by A_i (resp. B_i) the action space of Player MIN (resp. MAX) in state i ; $r_i^{ab} \in \mathbb{R}$ the transition payment, paid by Player MIN to Player MAX when the current state is i and the last actions of the players are $a \in A_i$ and $b \in B_i$, respectively; $P_i^{ab} \in \Delta(S)$ the transition probability from state i to the next state, assuming that the actions $a \in A_i$ and $b \in B_i$ were just chosen.

Starting from a given initial state i_0 , a zero-sum game with perfect information is played in stages as follows: at step ℓ , if the current state is i_ℓ , Player MIN chooses an action $a_{i_\ell} \in A_{i_\ell}$, and Player MAX subsequently chooses an action $b_{i_\ell} \in B_{i_\ell}$. Then, Player MIN pays $r_{i_\ell}^{a_{i_\ell} b_{i_\ell}}$ to Player MAX and the next state is chosen according to the probability law $P_{i_\ell}^{a_{i_\ell} b_{i_\ell}}$. By perfect information, we mean that both players are informed of the current state and of the previous actions of the other player. In particular, we assumed that Player MIN selects each action before Player MAX, who is therefore informed of this action.

A strategy σ of Player MIN (resp. τ of Player MAX) is a selection rule determining the current action of Player MIN (resp. MAX) as a function of the current and past information available to him (resp. her). Given strategies σ and τ of the two players, the *payoff* of the game in horizon k is the

following additive function of the transition payments:

$$J_i^k(\sigma, \tau) = \mathbb{E}_{i, \sigma, \tau} \left[\sum_{\ell=0}^{k-1} r_{i_\ell}^{a_\ell b_\ell} \right],$$

where $\mathbb{E}_{i, \sigma, \tau}$ denotes the expectation for the probability law of the process $(i_\ell, a_\ell, b_\ell)_{\ell \geq 0}$, determined by the transition probabilities, the strategies σ and τ of the players, and the initial state $i_0 = i$. Player MIN intends to choose a strategy minimizing the payoff J_i^k , whereas Player MAX intends to maximize the same payoff. The value of the k -stage game (played in horizon k) starting at state i is thus defined as

$$v_i^k = \inf_{\sigma} \sup_{\tau} J_i^k(\sigma, \tau),$$

the infimum and the supremum being taken over the set of strategies of the players MIN and MAX, respectively. Note that here the infimum and the supremum commute, since the information is perfect.

The study of the value vector $v^k = (v_i^k)_{i \in S}$ involves the *dynamic programming operator*, or *Shapley operator*. The latter is a map $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ whose i th coordinate is given by

$$T_i(x) = \inf_{a \in A_i} \sup_{b \in B_i} (r_i^{ab} + P_i^{ab} x), \quad x \in \mathbb{R}^n. \quad (1)$$

Note that an element $P \in \Delta(S)$ is seen as a row vector $P = (P_j)_{j \in S}$ of \mathbb{R}^n , so that Px means $\sum_{j \in S} P_j x_j$. The Shapley operator allows one to determine recursively the value vector: $v^k = T(v^{k-1})$ and $v^0 = 0$. See [16], [17], [18] for more background.

Here, we are interested in the asymptotic behavior of the sequence of mean values per time unit (v^k/k) as $k \rightarrow \infty$. When the latter ratio converges, the limit, denoted by $\chi(T)$, is called the *mean payoff vector*. Note that, according to the dynamic programming principle, the mean payoff vector is given by

$$\chi(T) = \lim_{k \rightarrow \infty} \frac{T^k(0)}{k},$$

where $T^k := T \circ \dots \circ T$ denotes the k th iterate of T .

The limit does exist for important classes of games. This question was studied in now classical papers of Bewley and Kohlberg [19] and Mertens and Neyman [20]. See also Rosenberg and Sorin [21], Sorin [22], Renault [23], and Bolte, Gaubert and Vigeral [24] for more recent developments. The existence of $\chi(T)$ is not guaranteed in general: a recent result of Vigeral [25] shows that the limit may not exist even if the action spaces are compact and the transition payments and probabilities are continuous functions of the actions.

B. Shapley operators

Given a zero-sum stochastic game with finite state space, its Shapley operator T , represented in (1), satisfies the following two properties:

- additive homogeneity: $T(x + \lambda e) = T(x) + \lambda e$, $\lambda \in \mathbb{R}$;
- monotonicity: $x \leq y \implies T(x) \leq T(y)$;

where \mathbb{R}^n is endowed with its usual partial order, and e is the *unit vector* of \mathbb{R}^n , i.e., the vector with all its entries equal to 1. The importance of these axioms in stochastic control and game theory has been known for a long time [26].

Kolokoltsov showed in [4] that, conversely, every operator from \mathbb{R}^n to \mathbb{R}^n that satisfies the latter two properties can be written as a Shapley operator (1). Rubinov and Singer [27] showed that the game may even be required to be deterministic. This motivates the following definition.

Definition 1: We call *Shapley operator* over \mathbb{R}^n any operator from \mathbb{R}^n to \mathbb{R}^n that is both monotone and additively homogeneous.

C. The ergodic eigenproblem

Given a Shapley operator T over \mathbb{R}^n , a basic question is to understand when the mean payoff vector is independent of the initial state. In other words, we ask whether $[\chi(T)]_i = [\chi(T)]_j$ holds for all $i, j \in S$, i.e., $\chi(T)$ is proportional to the *unit vector* e of \mathbb{R}^n .

This question is related to the nonlinear spectral problem consisting in finding an eigenpair $(\lambda, u) \in \mathbb{R} \times \mathbb{R}^n$ solution of the *ergodic equation*

$$T(u) = \lambda e + u. \quad (2)$$

Indeed, when the ergodic eigenproblem (2) has a solution, it is easy to see that $\chi(T) = \lambda e$ (this follows for instance from Proposition 2 below). Hence, the scalar λ , called the *eigenvalue* of T , yields the mean payoff per time unit for any initial state. Moreover, the vector u , called *bias vector* or *eigenvector*, gives optimal, or ε -optimal, stationary strategies by identifying in (1) the actions for which the infimum (resp. the supremum) is attained, or approached with an arbitrarily small precision.

D. Boundedness of slice spaces in Hilbert's seminorm

Our approach of the ergodic eigenproblem, following [13], relies on the study of the *sub-eigenspaces* $\mathcal{S}_\alpha(T)$, *super-eigenspaces* $\mathcal{S}^\beta(T)$ and *slice spaces* $\mathcal{S}_\alpha^\beta(T)$, for $\alpha, \beta \in \mathbb{R}$, which are defined respectively by:

$$\begin{aligned} \mathcal{S}_\alpha(T) &= \{x \in \mathbb{R}^n \mid T(x) \geq \alpha e + x\}, \\ \mathcal{S}^\beta(T) &= \{x \in \mathbb{R}^n \mid T(x) \leq \beta e + x\}, \\ \mathcal{S}_\alpha^\beta(T) &= \mathcal{S}_\alpha(T) \cap \mathcal{S}^\beta(T). \end{aligned} \quad (3)$$

These subsets are of intrinsic interest. Indeed, the following standard proposition shows that by checking the nonemptiness of any of these sets, we can bound the mean payoff from above, from below, or in both directions.

Proposition 2 ([22, Prop. 7]): Assume that $\mathcal{S}_\alpha(T) \neq \emptyset$. Then, $\liminf_{k \rightarrow \infty} T^k(0)/k \geq \alpha e$. Similarly, if $\mathcal{S}^\beta(T) \neq \emptyset$, then, $\limsup_{k \rightarrow \infty} T^k(0)/k \leq \beta e$.

Therefore, an element u such that $T(u) \geq \alpha e + u$ may be thought of as a *certificate* that $\liminf_{k \rightarrow \infty} T^k(0)/k \geq \alpha e$. A dual interpretation holds when $T(u) \leq \beta e + u$. Also, if the ergodic equation (2) has a solution, we deduce that $\mathcal{S}_\alpha^\lambda \neq \emptyset$, and so, it follows from Proposition 2 that $\chi(T) = \lambda e$.

Let us now introduce the *Hilbert's seminorm* on \mathbb{R}^n , which is given by $\|x\|_H = \max_{i \in S} x_i - \min_{i \in S} x_i$. It is a

useful tool in nonlinear Perron-Frobenius theory [12], [13], where it is sometimes called *Hopf's oscillation* [28], [29]. It also appears in linear consensus theory under the name of *diameter* or *Tsitsiklis' Lyapunov function* [30]. In this context, the set $\mathbb{R}e$ of scalar multiples of the unit vector is known as *consensus states*, and the Hilbert's seminorm, which vanishes on $\mathbb{R}e$, is a measure of the distance to consensus states. See [31], [32] for further developments concerning Hilbert's seminorm and the related Hilbert's projective metric in relation with consensus theory.

It is a standard result (see e.g. [13]) that any Shapley operator $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is nonexpansive with respect to Hilbert's seminorm, meaning that, for every $x, y \in \mathbb{R}^n$, $\|T(x) - T(y)\|_H \leq \|x - y\|_H$. Thus, as a special case of Theorem 4.1 in [12], we get that the ergodic eigenproblem (2) has a solution if, and only if, one orbit of T , $\{T^k(x) \mid k \in \mathbb{N}\}$ with $x \in \mathbb{R}^n$, is bounded in Hilbert's seminorm. Hence, the solvability of (2) boils down to finding a nonempty invariant set under T that is bounded in Hilbert's seminorm. It is readily seen that all the subsets in (3) are invariant under T . Then, we have the following result.

Theorem 3 ([12, Th. 4.1], [13]): Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a Shapley operator. Assume that there is a choice of $\alpha, \beta \in \mathbb{R}$ such that the slice space $\mathcal{S}_\alpha^\beta(T)$ is nonempty and bounded in Hilbert's seminorm. Then, there exist $\lambda \in \mathbb{R}$ and $u \in \mathbb{R}^n$ such that $T(u) = \lambda e + u$.

Note that the slice space $\mathcal{S}_\alpha^\beta(T)$ is nonempty as soon as $\alpha \leq \min_{i \in S} T_i(0)$ and $\beta \geq \max_{i \in S} T_i(0)$, since in this case $0 \in \mathcal{S}_\alpha^\beta$. Hence, the main difficulty when applying Theorem 3 is to check that the boundedness condition does hold for some specific α, β . We shall see in the next section that the boundedness of the slices spaces *for all* values of α, β has a combinatorial characterization. Also, we easily deduce the following result.

Corollary 4: Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a Shapley operator. If all the slice spaces $\mathcal{S}_\alpha^\beta(T)$ are bounded in Hilbert's seminorm, then, for every vector $g \in \mathbb{R}^n$, there is a scalar $\lambda \in \mathbb{R}$ and a vector $u \in \mathbb{R}^n$ such that $g + T(u) = \lambda e + u$.

III. HYPERGRAPH CONDITIONS FOR THE SOLVABILITY OF THE ERGODIC EIGENPROBLEM

A. Hypergraphs associated with a Shapley operator

A *directed hypergraph* is a pair (N, D) , where N is a set of *nodes* and D is a set of (directed) *hyperarcs*. A hyperarc d is an ordered pair $(t(d), h(d))$ of disjoint nonempty subsets of nodes; $t(d)$ is the *tail* of d and $h(d)$ is its *head*. We shall often write for brevity t and h instead of $t(d)$ and $h(d)$, respectively. When t and h are both of cardinality one, the hyperarc is said to be an *arc*, and when every hyperarc is an arc, the directed hypergraph becomes a directed graph. In the following, the term *hypergraph* will always refer to a directed hypergraph. For background on hypergraphs, we refer the reader to [33] and the references therein, and in particular to [34] for reachability problems.

To express the “behavior at infinity” of a Shapley operator $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$, we introduce a pair of hypergraphs associated with T , and denoted by $(\mathcal{H}^+, \mathcal{H}^-)$. Before giving their

construction, let us fix some notation. We still denote by S the set $\{1, \dots, n\}$, and if I is a subset of S , we denote by e_I the indicator vector of I , i.e., the vector of \mathbb{R}^n such that $[e_I]_i = 1$ if $i \in I$ and $[e_I]_i = 0$ if $i \in S \setminus I$. If $I = \{i\}$, we will write e_i instead of $e_{\{i\}}$.

We now explain the construction of the hypergraphs \mathcal{H}^+ and \mathcal{H}^- associated with T :

- the set of nodes for both \mathcal{H}^+ and \mathcal{H}^- is S ;
- the hyperarcs of \mathcal{H}^+ are the pairs $(J, \{i\})$ such that $\lim_{\alpha \rightarrow +\infty} T_i(\alpha e_J) = +\infty$;
- the hyperarcs of \mathcal{H}^- are the pairs $(J, \{i\})$ such that $\lim_{\alpha \rightarrow -\infty} T_i(\alpha e_J) = -\infty$.

These hypergraphs have an intuitive game theoretic interpretation when the action spaces A_i and B_i are finite. Then, it can be readily checked that there is a hyperarc $(J, \{i\})$ in \mathcal{H}^- if, and only if, Player MIN can choose an action $a \in A_i$ in such a way that for all actions $b \in B_i$ of Player MAX, $[P_i^{ab}]_j > 0$ holds for at least one state $j \in J$. In other words, such a hyperarc encodes the fact that starting from state i , Player MIN can force the next state to belong to J , with positive probability. Similarly, there is a hyperarc $(J, \{i\})$ in \mathcal{H}^+ if, and only if, for every action $a \in A_i$, Player MAX can choose an action $b \in B_i$ (depending on a) such that $[P_i^{ab}]_j > 0$ holds for at least one state $j \in J$. Therefore, starting from state i , Player MAX can force the next state to belong to J , with positive probability.

B. Reachability conditions for bounded slice spaces

We first recall the definition of reachability in hypergraphs. Let $\mathcal{G} = (N, D)$ be a hypergraph. A *hyperpath* of length p from a set of nodes $I \subset N$ to a node $j \in N$ is a sequence of p hyperarcs $(\mathbf{t}_1, \mathbf{h}_1), \dots, (\mathbf{t}_p, \mathbf{h}_p)$, such that $\mathbf{t}_i \subset \bigcup_{k=0}^{i-1} \mathbf{h}_k$ for all $i = 1, \dots, p+1$ with the convention $\mathbf{h}_0 = I$ and $\mathbf{t}_{p+1} = \{j\}$. We say that a node $j \in N$ is *reachable* from a subset I of N if there exists a hyperpath from I to j . Alternatively, the relation of reachability can be defined in a recursive way: j is reachable from I if either $j \in I$ or there exists a hyperarc (\mathbf{t}, \mathbf{h}) such that $j \in \mathbf{h}$ and every node of \mathbf{t} is reachable from the set I . A subset J of N is said to be *reachable* from a subset I of N if every node of J is reachable from I . We denote by $\text{reach}(I, \mathcal{G})$ the set of reachable nodes from I in \mathcal{G} . A subset I of N is *invariant* in the hypergraph \mathcal{G} if it contains every node that is reachable from itself, i.e., $\text{reach}(I, \mathcal{G}) \subset I$ (hence $\text{reach}(I, \mathcal{G}) = I$, since the other inclusion always holds). One readily checks that, for $J \subset N$, $\text{reach}(J, \mathcal{G})$ is the smallest invariant set in \mathcal{G} containing J .

We also need the following definition.

Definition 5: A pair (I, J) of subsets of S is said to be a pair of *conjugate subsets* with respect to the hypergraphs $(\mathcal{H}^+, \mathcal{H}^-)$ if $I \cap J = \emptyset$, $\text{reach}(J, \mathcal{H}^+) = S \setminus I$ and $\text{reach}(I, \mathcal{H}^-) = S \setminus J$. Such a pair is said to be *trivial* if one of the sets I, J is empty.

The following is our main result.

Theorem 6: Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a Shapley operator. All the slice spaces $S_\alpha^\beta(T)$ are bounded in Hilbert's seminorm if,

and only if, there are only trivial pairs of conjugate subsets with respect to the hypergraphs $(\mathcal{H}^+, \mathcal{H}^-)$.

The proof will be given in the extended version of the present conference article. Let us just mention that the “if” part is proved by combinatorial means, and the “only if” part by a topological degree theory argument.

Combining this theorem and Corollary 4, we get the following existence result for solutions of the ergodic eigenproblem (2).

Corollary 7: Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a Shapley operator. Assume that there are only trivial pairs of conjugate subsets with respect to the hypergraphs $(\mathcal{H}^+, \mathcal{H}^-)$. Then, for every vector $g \in \mathbb{R}^n$, there is a scalar $\lambda \in \mathbb{R}$ and a vector $u \in \mathbb{R}^n$ such that $g + T(u) = \lambda e + u$.

It has already been mentioned in the introduction (see also Subsection IV-C), that the hypergraph conditions of Corollary 7 reduce to the ones in [8] when the payments are bounded. Thus, the main result of [8] implies that the converse property in Corollary 7 does hold if, for every state $i \in S$, the payment function $(a, b) \mapsto r_i^{ab}$ is bounded. In a forthcoming work, we show that the converse property is in fact true with arbitrary payments.

C. Algorithmic aspects

Given a Shapley operator T over \mathbb{R}^n , the basic issue under consideration is to check whether the ergodic eigenproblem (2) is solvable for all operators $g + T$ with $g \in \mathbb{R}^n$.

Corollary 7 provides a combinatorial condition for this property to hold. This condition can be effectively checked as soon as the limits $\lim_{\alpha \rightarrow \pm\infty} T_i(\alpha e_J)$ arising in the definition of the hyperarcs of \mathcal{H}^\pm can be computed. This is generally the case in practice, as the Shapley operator is typically given explicitly in terms of actions, payments, and transition probabilities. To give an elementary example, we already noted at the end of Subsection III-A that if all the action spaces A_i, B_i in (1) are *finite*, then, checking whether there is a hyperarc $(J, \{i\})$ in \mathcal{H}^\pm reduces to elementary tests. Indeed, the discussion in that subsection implies that $(J, \{i\})$ is a hyperarc of \mathcal{H}^- (resp. \mathcal{H}^+) if, and only if

$$\max_{a \in A_i} \min_{b \in B_i} P_i^{ab} e_J > 0 \quad (\text{resp. } \min_{a \in A_i} \max_{b \in B_i} P_i^{ab} e_J > 0) .$$

Therefore, checking whether the hyperarc (J, i) belongs to \mathcal{H}^\pm can be done in $O(n|A_i||B_i|)$ elementary operations, where $|\cdot|$ denotes the cardinality of a set. However, the limits $\lim_{\alpha \rightarrow \pm\infty} T_i(\alpha e_J)$ are computable in more general situations (see Section IV-D), and so the interest of Corollary 7 is to lead to an effective condition in such cases.

Let us now recall how reachability conditions can be checked algorithmically once the hypergraphs have been computed. Indeed, the conditions of Corollary 7 boil down to check that, for every $I \subset S$, either $\text{reach}(S \setminus I, \mathcal{H}^+) \neq S \setminus I$ or $\text{reach}(I, \mathcal{H}^-) = S$.

It is known that in a hypergraph $\mathcal{G} = (N, D)$, the set of reachable nodes from a subset I can be computed in $O(\text{size}(\mathcal{G}))$ time, where $\text{size}(\mathcal{G})$ denotes the *size* of \mathcal{G} , and is equal to $|N| + \sum_{d \in D} |\mathbf{t}(d)| + |\mathbf{h}(d)|$, see [34]. Hence, the conditions of Corollary 7 can be checked in

time $O(2^n \text{size}(\mathcal{H}^\pm))$, counting for one time unit a call to an oracle computing a limit $\lim_{\alpha \rightarrow \pm\infty} T_i(\alpha e_j)$. However, $\text{size}(\mathcal{H}^\pm)$ may be exponential in the number of states, since the number of hyperarcs is only bounded by $n2^n$. The bound $O(2^n \text{size}(\mathcal{H}^\pm))$ may appear to be coarse at first sight, but it cannot be reduced to a polynomial bound unless $P = \text{coNP}$. Indeed, the restricted version of this problem concerning deterministic operators with finite action spaces reduces to the nonexistence of a nontrivial fixed point of a monotone Boolean function, a problem which has been shown to be coNP -hard by Yang and Zhao [10].

We refer the reader to [8] for more information on complexity issues, and refinements in the case of finite actions spaces (in this case, one arrives at better complexity bounds, still exponential in n). We also point out that there are simpler sufficient conditions, involving directed graphs instead of a hypergraph [13]. Such conditions can be checked using only a polynomial number of elementary operations (still counting for one unit the call to the oracle computing a limit), but they are less accurate.

IV. COMPARISON WITH EARLIER RESULTS AND ILLUSTRATIVE EXAMPLE

We first show that several earlier results follow as special cases of Theorem 6. Then, we give a simple example in which Theorem 6 allows one to show the solvability of the ergodic problem, whereas earlier conditions fail.

A. The ergodic eigenproblem and the recession operator

To address the problem of the solvability of the ergodic equation (2), Gaubert and Gunawardena introduced in [13] the *recession operator* associated with a Shapley operator T over \mathbb{R}^n . This operator is defined, for $x \in \mathbb{R}^n$, by $\hat{T}(x) = \lim_{\alpha \rightarrow +\infty} T(\alpha x)/\alpha$. Its existence is not always guaranteed, but it does exist for any Shapley operator represented as in (1) with bounded payments. Moreover, when it exists, it inherits from T the additive homogeneity and monotonicity properties. Furthermore, it is positively homogeneous, meaning that $\hat{T}(\alpha x) = \alpha \hat{T}(x)$ for every $\alpha \geq 0$. As a consequence, any vector proportional to the unit vector of \mathbb{R}^n is a fixed point of \hat{T} . We shall call such fixed points *trivial* fixed points.

Theorem 13 of [13] shows that if a Shapley operator T has a recession function \hat{T} whose only fixed points are trivial, then all the slice spaces of T are bounded in Hilbert's seminorm. An example given in the same paper shows that this sufficient condition is not necessary. Therefore, the present Theorem 6 solves the issue raised there of giving a necessary and sufficient condition.

B. Weakly convex maps and risk sensitive control

Cavazos-Cadena and Hernández-Hernández gave in [7], under a weak convexity property, a necessary and sufficient condition, in terms of graph, under which the ergodic eigenproblem (2) is solvable for all the Shapley operators $g + T$ with $g \in \mathbb{R}^n$. They also showed that this condition is equivalent to \hat{T} having only trivial fixed points.

The weak convexity property was motivated by applications to risk sensitive control. A typical example of weakly convex operator is the following Shapley operator T :

$$T_i(x) = \log \left(\sum_{j \in S} M_{ij} \exp(x_j) \right), \quad i \in S, \quad (4)$$

where $M = (M_{ij})$ is a nonnegative matrix without zero row. A supremum of weakly convex Shapley operators is weakly convex, hence, one can construct further examples of weakly convex operators by taking suprema of operators of the form (4).

It can be easily checked that under the weakly convex hypothesis, our hypergraph conditions (Theorem 6) reduce to the graph conditions of [7]. However, this weak convexity assumption is typically not satisfied for two-player Shapley operators, or for dynamic programming operators of general stochastic control problems.

C. Hypergraphs associated with a recession operator

A recent result of the authors [8] concerns the special case in which the payment functions $(a, b) \mapsto r_i^{ab}$ are *bounded* for all $i \in S$. Then, the recession map of the Shapley operator T exists and each of its coordinate functions is given by

$$\hat{T}_i(x) = \min_{a \in A_i} \max_{b \in B_i} P_i^{ab} x, \quad i \in S.$$

Theorem 3.1 of [8] shows that the ergodic equation is solvable for all operators $g + T$ with $g \in \mathbb{R}^n$ if, and only if, the recession operator \hat{T} has only trivial fixed points. It is also shown there that the latter condition can be characterized in terms of reachability in directed hypergraphs, depending only on the supports of the probability laws P_i^{ab} . We leave it as an exercise to check that the hypergraph conditions of Theorem 6 reduces to the one of [8] in this special case. It follows that the sufficient condition of Corollary 7 is necessary when the payment functions are bounded.

D. Example

Here, we illustrate the situation in which the conditions for the solvability of the eigenproblem (2) given by Corollary 7 are satisfied, whereas the conditions provided by [13], [7], [8] do not apply. Consider the parametric family of Shapley operators $g + T$, $g \in \mathbb{R}^3$, with $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by

$$T(x) = \begin{pmatrix} \sup_{0 < p \leq 1} (\log p + p(x_2 \wedge x_3) + (1-p)x_1) \\ \inf_{0 < p \leq 1} (-\log p + px_3 + (1-p)x_1) \\ x_3 \end{pmatrix}$$

where \wedge stands for min. The operator $g + T$ corresponds to a game with three states: Player MAX partially control state 1, Player MIN controls state 2 and state 3 is an absorbing state. In state 1, Player MAX first receives g_1 from Player MIN, then chooses an action $p \in (0, 1]$, and receives in addition $\log p$. Then, with probability $1 - p$, the next state remains 1, and with probability p , it is chosen by Player MIN between state 2 and state 3. Thus, maximizing the one day payoff would lead to select $p = 1$, but this leads to leave state 1 with probability one. A dual interpretation applies to Player MIN in state 2.

Figure 1 shows a concise representation of the hypergraphs \mathcal{H}^+ and \mathcal{H}^- associated with T , in which only the hyperarcs with minimal tail (with respect to the inclusion partial order) have been represented. To construct these hypergraphs, it is convenient to notice that $T_1(x) = h((x_2 \wedge x_3) - x_1) + x_1$ and $T_2(x) = -h(x_1 - x_3) + x_1$, where h is the real function defined by $h(z) = \sup_{0 < p \leq 1} (\log p + pz)$. Also note that h satisfies $h(z) = -1 - \log(-z)$ for $z \leq -1$, and $h(z) = z$ for $z \geq -1$. Thus, for instance, there is no arc from $\{2\}$ to $\{1\}$ in \mathcal{H}^+ since $T_1(\alpha e_2) = 0$ for all $\alpha \geq 0$. However, there is a hyperarc from $\{2, 3\}$ to $\{1\}$, since $T_1(\alpha e_{\{2,3\}}) = \alpha$ for all $\alpha \geq 0$, which yields $\lim_{\alpha \rightarrow +\infty} T_1(\alpha e_{\{2,3\}}) = +\infty$. Alternatively, following the discussion at the end of Subsection III-A, one may checked that, starting from state 1, Player MAX can force the next state to belong to $\{2, 3\}$ with positive probability, but he cannot force it to be 2.



Fig. 1. The hypergraphs associated with T

Then, one may check that there are no nonempty disjoint subsets of states I, J such that $\text{reach}(J, \mathcal{H}^+) = \{1, 2, 3\} \setminus I$ and $\text{reach}(I, \mathcal{H}^-) = \{1, 2, 3\} \setminus J$. Hence, there do not exist nontrivial conjugate subsets with respect to $(\mathcal{H}^+, \mathcal{H}^-)$. By application of Theorem 6, we deduce that all the slice spaces of T are bounded in Hilbert's seminorm. It follows that the ergodic equation (2) is solvable for all operators $g + T$, with $g \in \mathbb{R}^3$.

This conclusion cannot be obtained from the previous theorems in [13], [7], [8]. Indeed, although the recession operator of T exists (it is given by $\hat{T}_1(x) = x_1 \vee (x_2 \wedge x_3)$, $\hat{T}_2(x) = x_1 \wedge x_3$ and $\hat{T}_3(x) = x_3$), one may check that any vector $[\alpha, 0, 0]^\top$ with $\alpha \geq 0$ is a fixed point of \hat{T} . Thus, ergodicity conditions requiring the nonexistence of nontrivial fixed points of the recession operator \hat{T} cannot be applied.

REFERENCES

- [1] P. Whittle, *Optimization over time. Vol. II*, ser. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons, Ltd., Chichester, 1983.
- [2] O. Hernández-Lerma and J. B. Lasserre, *Further topics on discrete-time Markov control processes*, ser. Applications of Mathematics (New York). Springer-Verlag, New York, 1999, vol. 42.
- [3] J. Bather, "Optimal decision procedures for finite Markov chains. II. Communicating systems," *Advances in Appl. Probability*, vol. 5, pp. 521–540, 1973.
- [4] V. N. Kolokoltsov, "On linear, additive, and homogeneous operators in idempotent analysis," in *Idempotent analysis*, ser. Adv. Soviet Math. Amer. Math. Soc., Providence, RI, 1992, vol. 13, pp. 87–101.
- [5] V. N. Kolokoltsov and V. P. Maslov, *Idempotent analysis and its applications*, ser. Mathematics and its Applications. Kluwer Academic Publishers Group, Dordrecht, 1997, vol. 401.
- [6] R. Cavazos-Cadena and D. Hernández-Hernández, "Necessary and sufficient conditions for a solution to the risk-sensitive Poisson equation on a finite state space," *Systems Control Lett.*, vol. 58, no. 4, pp. 254–258, 2009.

- [7] —, "Poisson equations associated with a homogeneous and monotone function: necessary and sufficient conditions for a solution in a weakly convex case," *Nonlinear Anal.*, vol. 72, no. 7-8, pp. 3303–3313, 2010.
- [8] M. Akian, S. Gaubert, and A. Hochart, "Ergodicity conditions for zero-sum games," *Discrete Contin. Dyn. Syst.*, vol. 35, no. 9, pp. 3901–3931, 2015.
- [9] G. Olsder, "Eigenvalues of dynamic max-min systems," *Discrete Event Dynamic Systems*, vol. 1, no. 2, pp. 177–207, 1991.
- [10] K. Yang and Q. Zhao, "The balance problem of min-max systems is co-NP hard," *Systems Control Lett.*, vol. 53, no. 3-4, pp. 303–310, 2004.
- [11] Y. Oshime, "An extension of Morishima's nonlinear Perron-Frobenius theorem," *J. Math. Kyoto Univ.*, vol. 23, no. 4, pp. 803–830, 1983.
- [12] R. D. Nussbaum, "Hilbert's projective metric and iterated nonlinear maps," *Mem. Amer. Math. Soc.*, vol. 75, no. 391, pp. iv+137, 1988.
- [13] S. Gaubert and J. Gunawardena, "The Perron-Frobenius theorem for homogeneous, monotone functions," *Trans. Amer. Math. Soc.*, vol. 356, no. 12, pp. 4931–4950 (electronic), 2004.
- [14] M. Arisawa, "Ergodic problem for the Hamilton-Jacobi-Bellman equation. I. Existence of the ergodic attractor," *Ann. Inst. H. Poincaré Anal. Non Linéaire*, vol. 14, no. 4, pp. 415–438, 1997.
- [15] P. Cardaliaguet, "Ergodicity of Hamilton-Jacobi equations with a non-coercive nonconvex Hamiltonian in $\mathbb{R}^2/\mathbb{Z}^2$," *Ann. Inst. H. Poincaré Anal. Non Linéaire*, vol. 27, no. 3, pp. 837–856, 2010.
- [16] J. Filar and K. Vrieze, *Competitive Markov decision processes*. Springer-Verlag, New York, 1997.
- [17] A. Neyman and S. Sorin, Eds., *Stochastic games and applications*, ser. NATO Science Series C: Mathematical and Physical Sciences, vol. 570. Kluwer Academic Publishers, Dordrecht, 2003.
- [18] J.-F. Mertens, S. Sorin, and S. Zamir, *Repeated games*. Cambridge University Press, 2014, vol. 55.
- [19] T. Bewley and E. Kohlberg, "The asymptotic theory of stochastic games," *Math. Oper. Res.*, vol. 1, no. 3, pp. 197–208, 1976.
- [20] J.-F. Mertens and A. Neyman, "Stochastic games," *Internat. J. Game Theory*, vol. 10, no. 2, pp. 53–66, 1981.
- [21] D. Rosenberg and S. Sorin, "An operator approach to zero-sum repeated games," *Israel J. Math.*, vol. 121, pp. 221–246, 2001.
- [22] S. Sorin, "Asymptotic properties of monotonic nonexpansive mappings," *Discrete Event Dyn. Syst.*, vol. 14, no. 1, pp. 109–122, 2004.
- [23] J. Renault, "Uniform value in dynamic programming," *J. Eur. Math. Soc. (JEMS)*, vol. 13, no. 2, pp. 309–330, 2011.
- [24] J. Bolte, S. Gaubert, and G. Vigeral, "Definable zero-sum stochastic games," *Math. Oper. Res.*, vol. 40, no. 1, pp. 171–191, 2015.
- [25] G. Vigeral, "A zero-sum stochastic game with compact action sets and no asymptotic value," *Dyn. Games Appl.*, vol. 3, no. 2, pp. 172–186, 2013.
- [26] M. G. Crandall and L. Tartar, "Some relations between nonexpansive and order preserving mappings," *Proc. Amer. Math. Soc.*, vol. 78, no. 3, pp. 385–390, 1980.
- [27] A. M. Rubinov and I. Singer, "Topical and sub-topical functions, downward sets and abstract convexity," *Optimization*, vol. 50, no. 5-6, pp. 307–351, 2001.
- [28] E. Hopf, "An inequality for positive linear integral operators," *J. Math. Mech.*, vol. 12, pp. 683–692, 1963.
- [29] P. J. Bushell, "Hilbert's metric and positive contraction mappings in a Banach space," *Arch. Rational Mech. Anal.*, vol. 52, pp. 330–338, 1973.
- [30] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Automat. Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [31] R. Sepulchre, A. Sarlette, and P. Rouchon, "Consensus in non-commutative spaces," in *Decision and Control (CDC), 2010 49th IEEE Conference on*, Dec 2010, pp. 6596–6601.
- [32] S. Gaubert and Z. Qu, "Markov operators on cones and non-commutative consensus," in *Control Conference (ECC), 2013 European*, July 2013, pp. 2693–2700.
- [33] X. Allamigeon, "On the complexity of strongly connected components in directed hypergraphs," *Algorithmica*, vol. 69, no. 2, pp. 335–369, 2014.
- [34] G. Gallo, G. Longo, S. Nguyen, and S. Pallottino, "Directed hypergraphs and applications," *Discrete Appl. Math.*, vol. 42, no. 2-3, pp. 177–201, 1993.